IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In re application of:	) Confirmation No.: 5951
	)
Mark Gray	) Examiner: Ramsey Refai
	)
Serial No.: 09/835,059	) Group Art Unit No.: 3627
	)
Filed on: April 13, 2001	)
	)

For: METHOD AND APPARATUS FOR DETERMINING INTERCONNECTIONS OF NETWORK DEVICES

Commissioner for Patents
P. O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF

Sir:

This Appeal Brief is submitted (a) in response to the Final Office Action mailed September 24, 2008 and (b) with a Notice of Appeal filed on the same day herewith.

This Appeal Brief reinstates an appeal after prosecution was reopened on October 17, 2007 with the mailing of a non-final Office Action. According to MPEP § 1204.01, Applicants are only required to pay the difference between the current fees and the amount previously paid.

I. REAL PARTY IN INTEREST

Sun Microsystems, Inc. is the real party in interest.

II. RELATED APPEALS AND INTERFERENCES

Appellants are unaware of any related appeals or interferences.

III. STATUS OF CLAIMS

Claims 1-8, 11-16, 20-29, 31-40, 42-43, 46-51, and 54-68 are pending in this application. Claims 6-8, 11-16, 27-29, 38-40, 42, 43, 46-51, 54-57, and 63-68 were withdrawn from consideration. Claims 1-5, 20-26, 31-37, and 58-62 were finally rejected in the Final Office Action mailed on September 24, 2008. Claims 9-10, 17-19, 30, 41, 44-45, and 52-53 were canceled during prosecution.

Claims 1-5, 20-26, 31-37, and 58-62 are the subject of this appeal.

IV. STATUS OF AMENDMENTS

No amendments were filed after the Final Office Action mailed on September 24, 2008.

V. SUMMARY OF CLAIMED SUBJECT MATTER

The present application contains independent Claims 1, 20, and 31. Claim 1 is a method claim. Claim 20 is a computer readable medium counterpart of method Claim 1. Claim 31 is an apparatus counterpart of method Claim 1.

Claims 1, 20, and 31 are generally directly to an approach for determining whether a connection exists among a plurality of network devices. It is often desirable to know the details of the interconnections among devices in a network in order to create, modify, operate, and maintain such a network. (Application, page 3, lines 1-2.) However, as networks become

increasingly large and complex, it becomes increasingly difficult for a technician to physically trace all of the connections to ensure that they have been established correctly. (Application, page 4, lines 1-2.) One example of such a large and complex network is a “server farm,” in which servers and other network devices are made available for use by a variety of clients or companies. (Application, page 3, lines 8-9.)

Prior to the invention, one approach for determining the interconnections among the devices of a network was to establish a rigid specification or definition that prescribes how each network device is connected to other network devices. For example, for a network device “A,” such a definition might specify that network device "A" is to be connected via a connection “B” to a port “C” on a network device “D”. (Application, page 4, lines 3-7.)

Once a type of rigid definition is established, the network may be constructed according to detailed specifications of how each network device is to be connected to the other network devices. The rigid definition approach may be effective if the number of network devices is small because the limited number of interconnections makes it feasible to manually verify each connection between the network devices and correct any problems or mistakes. (Application, page 4, lines 8-13.)

However, as the number of network devices, the types of network devices, the number of connections and the types of connections all increase, the effectiveness of the rigid definitional approach decreases. In particular, the more network devices and connections there are, the more difficult it is to test each connection and identify those connections that are misconnected or otherwise have problems that are to be resolved. (Application, page 4, lines 14-18.)

Furthermore, because of the inevitability of human error, merely having personnel repeatedly check the work done against the rigid network definition is often not effective in

identifying a sufficient number of the problems with the network configuration. Inevitably some problems will escape even the most careful inspection by qualified personnel. (Application, page 5, lines 3-7.)

Each of Claims 1, 6, 12, 16, 20, 27-29, 31, and 38-40 provides a technique that may reduce the opportunity for human error in determining the interconnections among network devices.

A. CLAIMS 1, 20, AND 31

Claim 1 features determining logical interconnections among network devices by changing the power state of one device and identifying an alteration at a second device in response to the change of power state of the first device. By changing the power state of one device to identify alterations at other devices as in the approach of Claim 1, existing physical connections can be determined and information representing the corresponding logical connection created and stored.

Claim 1 recites the power state of a first network device is changed from either (a) an unpowered state to a powered state or (b) from the powered state to the unpowered state. (Application, page 10, lines 20-25; FIG. 2, block 220.) For example, the power state of a device may be changed by causing a CPU that is initially unpowered (or "off") to have power supplied to the CPU (e.g., the CPU is turned on).

Claim 1 then recites "identifying whether an alteration occurs at a second network device in response to changing the power state of the first network device." (Application, page 11, lines 11-17; FIG. 2, block 230.) As a specific example, a control device can identify an alteration at a switch in response to turning a CPU "on," such as the raising of a trap on the

port of the switch, which would indicate that the CPU and switch are connected through the port on which the trap is raised. (Application, page 11, lines 18-23.)

Claim 1 finally recites "when the alteration occurs at the second network device, creating and storing first information representing a logical connection of the first network device to the second network device" (Application, page 12, lines 7-16; FIG. 2, blocks 240, 250.)

Independent Claim 20 is a computer-readable medium counterpart of method Claim 1, and includes limitations analogous to the limitations of Claim 1. Thus, the elements of Claim 20 are disclosed in at least the same sections of the Specification and Drawings as those cited above in connection with Claim 1. In addition, the elements of Claim 20 are supported by the hardware and computer-readable media description provided on page 48, line 17 through page 51, line 20.

Independent Claim 31 is an apparatus counterpart of method Claim 1, and includes limitations analogous to the limitations of Claim 1. Thus, the elements of Claim 31 are disclosed in at least the same sections of the Specification and Drawings as those cited above in connection with Claim 1. In addition, the elements of Claim 31 are supported by the hardware description provided on page 48, line 17 through page 51, line 20. Furthermore, the "means for changing the power state of the first network device" may be the power controller 120 of FIG. 1 (see also Application, page 11, lines 1-4); the "means for identifying whether an alteration occurs at a second network device in response to changing the power state of the first network device" may be the control device 110 of FIG. 1 (see also Application, page 11, lines 18-23); and the "means for creating and storing first information" may be the control device 110 and database 112 of FIG. 1 (see also Application, page 12, lines 7-10).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

Claims 1-3, 5, 20-22, 31-33, and 35 stand rejected under 35 U.S.C. § 103(a) as being allegedly unpatentable over U.S. Patent No. 6,728,670 issued to Schenkel et al. ("*Schenkel*") in view of U.S. Patent No. 6,516,345 issued to Kracht ("*Kracht*").

Claims 4, 23, and 34 stand rejected under 35 U.S.C. § 103(a) as being allegedly unpatentable over *Schenkel* in view of *Kracht* and in further view of U.S. Patent No. 6,628,623 issued to Noy ("*Noy*").

Claims 25, 26, 36, 37, 58, and 59 stand rejected under 35 U.S.C. § 103(a) as being allegedly unpatentable over *Schenkel* in view of *Kracht* and in further view of U.S. Patent No. 5,347,167 issued to Singh ("*Singh*").

Claims 60-62 stand rejected under 35 U.S.C. § 112(1) as allegedly not being supported in the Specification.

VII. ARGUMENTS

It is respectfully submitted that the Examiner has erred in rejecting Claims 1-5, 20-26, 31-37, and 58-62 under 35 U.S.C. §103(a).

A. CLAIMS 1-3, 5, 20-22, 31-33, AND 35

Claims 1-3, 5, 20-22, 31-33, and 35 stand rejected under 35 U.S.C. § 103(a) as being allegedly unpatentable over *Schenkel* in view of *Kracht.*

1. Claim 1

Claim 1 features:

> "A method for determining one or more logical interconnections among a plurality of network devices that are interconnected in a network in an indefinite relationship, wherein a power state is associated with a first network device, the method comprising the computer-implemented steps of:
>
> **changing the power state** of the first network device **from either (a) an *unpowered* state to a *powered* state or (b) from the *powered* state to the *unpowered* state**;
>
> **identifying whether an <u>alteration</u> occurs *at a second network* device in response to changing the power state *of the first network device***; and
>
> when the alteration occurs at the second network device, creating and storing first information representing a logical connection of the first network device to the second network device." (Emphasis added.)

i) The Final Office Action's Citations From *Schenkel*

The Final Office Action states that *Schenkel* discloses "changing the power state of the first network device from either (a) an unpowered state to a powered state or (b) from the powered state to the unpowered state; identifying whether an alteration occurs at a second network device in response to changing the power state of the first network device (column 2,

lines 20-40; shows a signal sent from a source device to a destination device, Figure 2, and column 3, lines 18-32.)" This is incorrect.

The first cited portion from Column 2 of *Schenkel* describes measuring the traffic output of one device (e.g., the sequence of bursts of packets formed of orthogonal signals), measuring the traffic input of another device, and determining connections between devices or a sequence of connections between devices based on whether the measured traffic between the two devices is statistically the same or not. (Col. 2, lines 20-40.) The last cited portion from Column 3 of *Schenkel* describe a series of four devices, A through D, connected in series in which the output of one device is the input to the next device, as illustrated in Figure 2. (Col. 3, lines 18-32; Figure 2.) Thus, the cited portions of *Schenkel* describe sending traffic from the source device to a destination device and comparing the traffic sent from the source device to the destination device. If the traffic is statistically the same, *Schenkel's* approach is to conclude that the source device is connected to the destination device, otherwise if the traffic is not statistically the same, the source device is not connected to the destination device.

The "Response to Arguments" section of the Final Office Action also cites additional portions of *Schenkel* – namely the Abstract, col. 2, lines 11-12; col. 1, line 65 – col. 2, line2; col. 22, line 60 – col. 23, line 15; and col. 19, lines 10-67. However, the cited portion of Column 19 of *Schenkel* describes and defines an "idle" device as a device in which the "traffic in or out of it is insignificant...Idleness can be expressed as having a mean level of traffic below some cutoff to be chosen by the operator." (Col. 19, lines 34-36 and 41-42.) Thus, because *Schenkel's* device is receiving traffic, the device must be powered, and when receiving more traffic so that the device is no longer idle, the device remains powered.

As described below, none of these portions of *Schenkel* describe either "**changing the power state** of the first network device from either (a) an ***unpowered state*** to a ***powered state***

or (b) from the ***powered state*** to the ***unpowered state***" or "identifying whether an alteration occurs at a second network device in response to changing the power state of the first network device" because (1) the sending of packets from a source device to a destination device in *Schenkel* does not change the power state of the destination device, and (2) even if it did, any alteration occurs at the destination device, not the source device. These two arguments are fully outlined in the following sections.

The Applicant notes that the Examiner explained during the Interview conducted on December 15, 2005 that a cited reference must be read as a whole, and therefore that the Applicant should not solely focus on the portions cited from *Schenkel*. The Applicant has reviewed not just the cited portions of *Schenkel*, but the entirety of *Schenkel*, yet the Applicant has failed to find anything that supports the rejections in the Final Office Action.

Furthermore, the Applicant notes that according to the MPEP, in an Office Action "the particular part relied on must be designated as nearly as practicable ... The pertinence of each reference, if not apparent, must be clearly explained ..." (MPEP §707, citing 37 C.F.R. §1.104(c)(2)), and "the particular figure(s) of the drawings(s), and/or page(s) or paragraph(s) of the reference(s), and/or any relevant comments briefly stated should be included." (MPEP §707). Thus, the Applicant respectfully requests that if the Examiner believes other portions of *Schenkel* not cited in the Final Office Action disclose the features of the claims that the Applicant has been unable to locate, that the Examiner provide citations to those portions of *Schenkel* along with an explanation as to why the Examiner believes the disclosure in those portions of *Schenkel* disclose the features of the claims.

ii) *Schenkel* Fails to Show Changing the **Power State** of a **First** Network Device and Identifying an **Alteration** at a **Second** Network Device

As discussed during the Interview with the Examiner, the Applicant is unclear about which portions of *Schenkel* are being relied upon as showing the following features of Claim 1: (a) "the first network device", (b) "the second network device", (c) "the power state of the first network device," (d) "the alteration occurs at the second network device." The Applicant's attempts at matching the devices and discussion of *Schenkel* to the first two features of Claim 1 (e.g., the first and second network devices) results in inconsistencies with two other features of Claim 1 (e.g., the power state and the alteration).

It initially appears that the Final Office Action is equating the "destination device" and "source device" of *Schenkel* to the "first network device" and "second network device," respectively (e.g., items (a) and (b) above), of Claim 1 because the Final Office Action says that the "signal bursts are sent to the destination device until no longer idle, which is a change of the power state." Assuming for the moment that sending the signal bursts is a change of power state, this matching of *Schenkel's* devices to those of Claim 1 is consistent with feature (c) above because the first network device (e.g., the destination device) has its power state being changed. However, this is inconsistent with feature (d) above of "identifying whether an alteration occurs at the second network device" because that would mean an alteration occurs at the source device in *Schenkel* that sends the signal bursts (or some other device not described).

Furthermore, in *Schenkel's* approach, the link between the source device and destination device is determined by a statistical comparison of the traffic at the destination device and the source device. Thus, a change that occurs, if a change does occur at all, is at the destination device, not the source device, which is the same device at which the power state changes. Yet in the approach of Claim 1, in response to the change in the power state of the first network

device, an alteration occurs at the second network device. Thus, the Applicant respectfully submits that based on this first application of the elements disclosed in *Schenkel* to the features of Claim 1, *Schenkel* fails to disclose “identifying whether an alteration occurs at the second network device in response to changing the power state of the first network device.”

Alternatively, if the source and destination device are reversed such that the “first network device” of Claim 1 is the “source device” of *Schenkel* and the “second network device” of Claim 1 is the “destination device” in *Schenkel*, then that would be consistent with feature (d) above in that an alteration is identified at the destination device (e.g., the change, if any, in the packets that are sent to the destination device). However, this is inconsistent with feature (c) above of “the power state of the first network device” because that would mean that the power state of the source device is changed. However, in *Schenkel's* approach, it is the power state of the destination device that is allegedly changed by sending the signal burst of packets, not the source device. Thus, the Applicant respectfully submits that based on this second application of the elements disclosed in *Schenkel* to the features of Claim 1, *Schenkel* fails to disclose “changing the power state of the first network device from either (a) an unpowered state to a powered state or (b) from the powered state to the unpowered state.”

To summarize, regardless of how the source and destination devices of *Schenkel* are matched against the first and second network devices of Claim 1, *Schenkel* always has **both** the change in power state **and** the alteration at the destination device, yet in Claim 1, the change in power state and the alteration occur at ***different*** network devices, namely the first and second network devices.

iii) *Schenkel* Identifies Connected Devices When the Traffic is the Same, Whereas Claim 1 Identifies Connected Devices When an Alteration Occurs

The link between the source and destination devices in *Schenkel* is only determined if the traffic is statistically the same, and if the traffic is not statistically the same, then there is no link between the source and destination devices. But when the traffic is the same, then **there is no alteration** (or difference) in the traffic between the two devices in *Schenkel*. Only if the traffic is ***<u>not</u>* statistically the same** is there not a link determined between the source and destination devices. In other words, *Schenkel's* approach only identifies that two devices are connected if there is **<u>no</u>** difference in the traffic, but if the traffic *<u>is different</u>*, then there is **<u>no</u>** link.

Yet the approach of Claim 1 is the opposite of that of *Schenkel*. Specially, Claim 1 expressly features that the logical connection is created and stored when there is a difference, i.e., when an alteration ***does occur*** at the second device. But if this situation occurs with *Schenkel's* approach, the opposite conclusion is reached, namely that the source and destination devices are ***<u>not</u>*** connected.

iv) *Schenkel's* "**Idle**" Device is Not an "**Unpowered**" Device

Contrary to the assertions of the Final Office Action, the mere sending of a signal comprised of a sequence of packet bursts is not the same as "changing the power state of the first network device **from either (a) an *unpowered* state to a *powered* state or (b) from the *powered* state to the *unpowered* state**" as in Claim 1. In *Schenkel*, the sending of packet bursts does not change the power state of the sending device, the receiving device, or any other device, which is a fundamental difference between *Schenkel* and the approach of Claim 1. In fact, some changes to the power state of a sending device, such as from powered to unpowered,

would render the approach of *Schenkel* inoperative because the sending device would be incapable of sending the signal. Even other power state changes, such as by turning a device from unpowered to powered, would not result in sending the sequence of bursts of packets as disclosed in *Schenkel.*

In the "Response to Arguments" section, the Final Office Action states that the "idled device in Schenkel **can be taken as being an unpowered device**, since the idled device **does not have enough traffic activity to be considered an active device** on the network. The device is then stimulated using signal burst to an active state to allow the device's connections to be identified directly" (page 3, Argument B; emphasis added).

The Applicant respectfully disagrees that sending signal bursts to an idle device until the device is no longer idle is a change of the power state, based on the definition of an "idle" device and "idleness" as provided in *Schenkel.* Specifically, *Schenkel* states:

> Stimulation of idle devices in a network allow their connections to be identified directly. The present invention can determine that a device is **<u>idle</u>** ***because the volume of traffic in or out of it is*** **<u>insignificant</u>**. It can then instruct a signal burst to be sent to or across this device in order to generate ***enough traffic*** to accurately locate it in the network...***<u>Idleness</u>*** can be expressed as having ***a mean level of traffic below some*** **<u>cutoff</u>** to be chosen by the operator. A convenient value of this cutoff is 5 units of activity per sampling period as this provides the classic chi-squared formulation with sufficient data for its basic assumptions to be reasonable accurate. (Col. 19, lines 33-46; emphasis added.)

Therefore, ***Schenkel*** **clearly defines an idle device as a device for which the traffic is not zero, but *merely <u>insignificant</u>***, meaning that the traffic through the device is below a cutoff value does not allow for accurate identification of the network connections. The use of a signal burst to increase the traffic for an idle device so that the device can be located indicates that the device is already in a "powered" power state (e.g., the device is "on"). The sending of the signal burst does not change the power state from unpowered to powered or from off to on (or

vice versa). Rather, the signal burst supplies sufficient traffic so that the statistical comparison of the traffic sent to the traffic received is meaningful. Because the basis for *Schenkel's* connection identification approach is a statistical method, sufficient traffic must be used in order to make a statistically meaningful comparison between the traffic sent and the traffic received, and therefore conclude that the sending device and the receiving device are connected.

Thus, **an "*idle*" device as defined in *Schenkel* is a device that is in the "*powered*" power state (or "on"),** as opposed to an "unpowered" power state (or "off"). Changing the status of the device in *Schenkel* from "idle" to "not idle" merely means that there is sufficient traffic through or to the device for a statistically meaningful comparison of traffic sent versus traffic received, but the power state of the device remains unchanged in the powered or "on" power state. **If the initial power state of an idle device were unpowered or off, then the device would be unable to receive the signal burst in *Schenkel*'s approach.**

In contrast to *Schenkel*, Claim 1 features "changing the power state of the first network device **from either (a) an *unpowered* state to a *powered* state or (b) from the *powered* state to the *unpowered* state**." Neither the cited portions of *Schenkel* or any other portion of *Schenkel* discloses anything about changing the power state from "an unpowered state to a powered state" as featured in Claim 1, because *Schenkel's* technique of changing the status of an "idle" device, that is already powered but merely has too little traffic to accurately use *Schenkel's* statistics-based connection identification approach, by sending a signal burst merely increases the traffic to the device still leaves the device in a powered power state. Furthermore, there is nothing in either the cited portions of *Schenkel* or any other portion of *Schenkel* about changing the power state from "the powered state to the unpowered state," as featured in Claim 1. Indeed, *Schenkel* does not even mention the terms "power" or "state."

v) The Final Office Action Improperly Relies on Reading into Claim 1 a Definition of a Term Not Used in Claim 1

During the Interview and in the Response to Arguments section of the Final Office Action, the Examiner explained that the Examiner was relying upon the definition of "unpowered" from page 10 the specification. However, the definition provided therein is **not** of the term "unpowered" but rather of the term "power cycling." Specifically, the specification states:

> Next, the power state of a device is changed, as indicated in block 220. For example, in FIG. 1, the initial power state of CPU 130 may be unpowered (or "off"), but then power is supplied to CPU 130 (e.g.,. CPU 130 is turned on). The changing of a power state may be referred to as "**power cycling**." However, **that term** is used herein in a broader sense to also include turning off a network device or even to change the power state of a network device from standby to active. (Application, page 10, lines 20-25; emphasis added.)

Note that in this portion of the Application, the term "power cycling" encompasses three types of power state changes: (1) from unpowered to powered, such as going from not being supplied with power to being supplied with power; (2) from "off" to "on;", and (3) from standby to active. This is consistent with other portions of the Application. For example, the Application describes "power cycling" as follows:

> The "**power cycling**" of a network device means that the power state of the network device is changed or altered from what the power state was immediately prior to the power cycling action. The power state of a network device before **power cycling** may simply be ***"off," unpowered, or inactive***, or ***"on," powered, or active***. The power state of a network device may also be any other power characteristic of the network device. For example, the power state may be a form of power conservation mode, such as a ***power saving or "sleep" state***, in which only minimal power is used by the network device. (Application, page 22, line 20 to page 23, line 2; emphasis added.)

The only place in the claims where the term "power cycling" is used is in Claims 16, 29, and 40 in which the term is expressly limited to one type of change of power state, namely

"power cycling a first network device from either 'off' to 'on' or from 'on' to 'off'." The term "power cycling" is not used in Claim 1 or in any of the other independent claims. Rather, Claim 1 expressly features "changing the power state of the first network device," and similar to Claims 16, 29, and 40, the type of change of power state is expressly limited by the words of Claim 1 to "from either (a) an unpowered state to a powered state or (b) from the powered state to the unpowered state." Thus, Claim 1 expressly excludes changing the power state from standby to active.

Furthermore, the Applicant respectfully disagrees with the Final Office Action's attempt to read into Claim 1 a definition of a term not used in Claim 1. If the Applicant wanted Claim 1 to recite the term "power cycling," the Applicant would have included the term in Claim 1 (just as the Applicant has included that term in Claim 16). Yet by expressly not using the term "power cycling" in Claim 1, the Applicant has expressly differentiated Claim 1 from the definition of the term "power cycling." In addition, the Applicant notes that while the specification defines the term "power cycling" as changing the power state of a device, the specification has not defined the concept of changing the power state in terms of the term "power cycling," and therefore the Applicant respectfully submits that it is not proper for the Final Office Action to do so.

The only way that the Applicant has been able to reconcile the Final Office Action's reliance on the definition of "power cycling" from page 10 of the Application would be to equate the term "standby" to "unpowered." Yet the two terms are clearly distinguished in the Application, as evidenced by the two passages provide above. Furthermore, on its face, the term "standby" would be understood by one of ordinary skill in the art to mean that a device in "standby" has power, and thus a change in power state from "standby" to "active" leaves the device in a "powered" power state.

While *Schenkel* discloses an approach for determining a network topology by sending a signal consisting of a sequence of bursts of packets and measuring such packet traffic at the output of a sending device and the input of a receiving device, including the stimulation of an "idle" device for which the traffic is too low to make an accurate statistical comparison, this does not relate to "changing the power state of the first network device **from either (a) an *unpowered* state to a *powered* state or (b) from the *powered* state to the *unpowered* state**" as featured in Claim 1 of the present application.

vi) Conclusion of Discussion of Claim 1

Because *Schenkel* fails to disclose, teach, suggest, or in any way render obvious either "changing the power state of the first network device from either (a) an unpowered state to a powered state or (b) from the powered state to the unpowered state" or "identifying whether an alteration occurs at the second network device in response to changing the power state of the first network device," the Applicant respectfully submits that, for at least the reasons stated above, Claim 1 is allowable over *Schenkel* and is in condition for allowance.

2. *Claims 20 and 31*

Claims 20 and 31 contain features that are either the same as or similar to those described above with respect to Claim 1. For example, Claims 20 and 31 both feature "changing the power state of the first network device **from either (a) an *unpowered* state to a *powered* state or (b) from the *powered* state to the *unpowered* state**," which is the same as in Claim 1.

Therefore, based on at least the reasons stated above with respect to Claim 1, the Applicant respectfully submits that Claims 20 and 31 are allowable over the art of record and are in condition for allowance.

5. Claims 2-3, 5, 21-24, and 32-35

Claims 2-3 and 5 are dependent upon Claim 1, Claims 21-24 are dependent upon Claim 20, and Claims 32-35 are dependent upon Claim 31.

Thus, each of Claims 2-3, 5, 21-24, and 32-35 include each and every feature of the corresponding independent claims. Therefore, the Applicant respectfully submits that each of Claims 2-3, 5, 21-24, and 32-35 is therefore allowable for the reasons given above for the Claims 1, 20, and 31. In addition, each of Claims 2-3, 5, 21-24, and 32-35 introduces one or more additional limitations that independently render it patentable. Some of these additional features of the dependent claims are addressed below, while a full discussion of each dependent claim is not included herein at this time based on the fundamental differences already identified herein.

i) Claims 3, 22, and 33

Claims 3, 22, and 33 each features "determining whether a state of a port of the terminal server is changed from dead to active in response to changing the power state of the first network device." As a preliminary matter, there appears to be a typographical omission in the citations for Claims 3, 22, and 33 in the Final Office Action that begins by referring to "column 30 –37" since it is unclear what column and line numbers are being referred to.

The Final Office Action cites Col. 2, line 65 – Col. 3, line 7, of *Schenkel* as disclosing a terminal server, yet the Applicant does not see a terminal server listed or described. The word "terminal" does not even appear in that passage, and the only occurrence of the word "server" is in referring to "file servers," which clearly are not terminal servers. An electronic search of *Schenkel* has failed to find any other reference to a "terminal server," nor has the Applicant been able to find any other type of device within *Schenkel* that functions as a terminal server.

Next, while Col. 6, lines 30-35, lines 55-56, and Col. 27, lines 55-62, of *Schenkel* all refer to a "port," there is nothing in those cited portions or any other that the Applicant has found about the port being part of a terminal server. Furthermore, there is nothing in those cited portions of *Schenkel* about the state of the port changing from dead to active as in Claims 3, 22, and 33, little less that such a change in state is in response to changing the power state of the first network device, as in Claim 1. While the last citation refers to "port level of activity," it is in the context of receiving a burst, which means non-zero activity, and thus does not disclose anything about the port being dead.

Because *Schenkel* fails to disclose, teach, suggest, or in any way render obvious either "determining whether a state of a port of the terminal server is changed from dead to active in response to changing the power state of the first network device," the Applicant respectfully submits that, for at least the reasons stated above, Claims 3, 22, and 33 is allowable over *Schenkel* and are in condition for allowance.

B. CLAIM 4, 23, AND 34

Claims 4, 23, and 34 stand rejected under 35 U.S.C. § 103(a) as being allegedly unpatentable over *Schenkel* in view of *Kracht* and in further view of *Noy*.

Claim 4 is dependent upon Claim 1, Claim 23 is dependent upon Claim 20, and Claim 34 is dependent upon Claim 31. Thus, each of Claims 4, 23, and 34 include each and every feature of the corresponding independent claims. Therefore, the Applicant respectfully submits that each of Claims 4, 23, and 34 is therefore allowable for at least the reasons given above for the Claims 1, 20, and 31.

C. CLAIMS 25, 26, 36, 37, 58, AND 59

Claims 25, 26, 36, 37, 58, and 59 have been rejected under 35 U.S.C. § 103(a) as being allegedly unpatentable over *Schenkel* in view of *Kracht* and in further view of *Singh.*

Claims 58-59 are dependent upon Claim 1, Claims 25-26 are dependent upon Claim 20, and Claims 36-37 are dependent upon Claim 31. Thus, each of Claims 25, 26, 36, 37, 58, and 59 includes each and every feature of the corresponding independent claims. Therefore, the Applicant respectfully submits that each of Claims 25, 26, 36, 37, 58, and 59 is therefore allowable for at least the reasons given above for Claims 1, 20, and 31.

D. CLAIMS 60-62

Claims 60-62 stand rejected under 35 U.S.C. § 112(1) as allegedly not being supported in the specification.

Claims 60-62 each feature that "when the power state of the first network device is the unpowered state, the first network device is not able to receive one or more packets over the network" and "when the power state of the first network device is the powered state, the first network device is able to receive one or more packets over the network."

In reference to Claims 60-62, the Final Office Action states, "The claim(s) contains subject matter, which was not described in the specification in such a way as to reasonably convey to one skilled in the relevant art that the inventor(s), at the time the application was filed, had possession of the claimed invention." However, MPEP § 2163.05 states that "each claim limitation must be expressly, **implicitly, or inherently** supported in the originally filed disclosure" (emphasis added).

It is respectfully submitted that Claims 60-62 are at least implicitly or inherently fully supported by the Application, and no new matter is included. For example, the two portions of

the Application provided and discussed above (e.g., under section VII(A)(5)) use the terms "unpowered" and "powered," and both terms are distinguished from other terms, such as "standby" and "active" as well as "on" and "off." Furthermore, the term "unpowered" on its face means having "a lack of power," while "powered" means having "power."

In the context of network devices, one of ordinary skill in the art would understand that a network device in the "unpowered" state would mean that the network device is not able to receive network traffic in the form of packets in a packet-based network such as the worldwide packet-based network known as the Internet. (Application, page 32, lines 5-6.) Conversely, one of ordinary skill in the art would understand that a network device in the "powered" state would mean that the network device is able to receive network traffic in the form of packets in a packet-based network such as the Internet.

The Final Office Action's rejections are based on *Schenkel's* disclosure of sending of a signal in the form of a burst of packets (or sequential bursts of packets) from a source device to a destination device, which the Final Office Action alleges is the same as changing the power state of a device from unpowered to powered. However, as defined in Claims 60-62, "the first network device being in the unpowered state means that the first network device is not able to receive one or more packets over the network." Thus, in the approach of Claims 60-62, the first network device cannot change from the unpowered state to the powered state based on sending a signal of a burst of packets because, by definition, the first device being unpowered means that the first device cannot receive packets. Thus, because Claims 60-62 directly contradict the approach of *Schenkel*, the Applicant respectfully submits that Claims 60-62 are allowable and are in condition for allowance.

E. CONCLUSION AND PRAYER FOR RELIEF

Based on the foregoing, it is respectfully submitted that the rejection of Claims 1-5, 20-26, 31-37, and 58-62 under 35 U.S.C. § 103(a) being unpatentable over the cited art lacks the requisite factual and legal bases. Appellants therefore respectfully request that the Honorable Board reverse the rejection of Claims 1-5, 20-26, 31-37, and 58-62 under 35 U.S.C. § 103(a).

Respectfully submitted,

HICKMAN PALERMO TRUONG & BECKER LLP

/DanielDLedesma#57181/
Daniel D. Ledesma
Reg. No. 57,181

Date: November 4, 2008

2055 Gateway Place, Suite 550
San Jose, CA 95110-1089
Telephone: (408) 414-1080 ext. 229
Facsimile: (408) 414-1076

VIII. CLAIMS APPENDIX

1. (Previously Presented) A method for determining one or more logical interconnections among a plurality of network devices that are interconnected in a network in an indefinite relationship, wherein a power state is associated with a first network device, the method comprising the computer-implemented steps of:

changing the power state of the first network device from either (a) an unpowered state to a powered state or (b) from the powered state to the unpowered state;

identifying whether an alteration occurs at a second network device in response to changing the power state of the first network device; and

when the alteration occurs at the second network device, creating and storing first information representing a logical connection of the first network device to the second network device.

2. (Original) The method as recited in Claim 1, further comprising the steps of:

retrieving second information from a database, wherein the second information represents one or more logical connections of the first network device to the second network device;

comparing the second information from the database with the first information; and

generating an error if the second information indicates that a logical connection exists between the first and second network devices but the first information does not indicate that the logical connection exists between the first and second network devices.

3. (Original) The method as recited in Claim 1, wherein the second network device is a terminal server and wherein the step of identifying whether the alteration occurs at the terminal server further comprises:

determining whether a state of a port of the terminal server is changed from dead to active in response to changing the power state of the first network device.

4. (Original) The method as recited in Claim 1, wherein the second network device is a switch and wherein the step of identifying whether the alteration occurs at the switch further comprises:

determining whether a trap on a port of the switch is raised in response to changing the power state of the first network device.

5. (Original) The method as recited in Claim 1, further comprising:

receiving, in response to changing the power state of the first network device, additional information from the first network device; and

recording the additional information.

6. (Withdrawn) A method for determining one or more logical interconnections among a set of specified network devices that are interconnected in a network in an indefinite relationship, the method comprising the steps of:

(1) establishing connections among a plurality of network devices based upon a set of rules;

(2) activating a particular network device of said set of specified network devices by supplying power to the particular network device that previously was not supplied with power;

(3) identifying whether, in response to activating the particular network device, a change occurs at one or more network devices of said plurality of network devices;

(4) when the change occurs at each of the one or more network devices, creating and storing information representing a logical connection of the particular network device to each of the one or more network devices; and

(5) repeating steps (2), (3), and (4) for each of said set of specified network devices.

7. (Withdrawn) The method as recited in Claim 6, wherein the set of rules are applied based upon one or more attributes of each connection.

8. (Withdrawn) The method as recited in Claim 7, wherein the one or more attributes of each connection include information that is selected from the group consisting of a type of connection between two or more network devices, the number of connections between a specific network device and one or more other network devices, and that a particular connection is between a first type of network device and a second type of network device.

9.-10. (Cancelled)

11. (Withdrawn) The method as recited in Claim 6, wherein the step of identifying whether the change occurs at one or more network devices further comprises:

determining whether a trap on a port of each of the one or more network devices is raised in response to activating the particular network device by supplying power to the particular network device that previously was not supplied with power.

12. (Withdrawn) A method for determining how devices are interconnected in a network, the method comprising the computer-implemented steps of:

sending a signal from a control device that results in a change in a power state of a first network device in response to the signal, wherein the power state changes from either powered to unpowered or from unpowered to powered;

determining whether the first network device is connected to a second network device by identifying an alteration at the second network device that occurs in response to changing the power state of the first network device; and

when the alteration occurs at the second network device, creating and storing information representing that the first network device is connected to the second network device.

13. (Withdrawn) The method as recited in Claim 12 wherein the first network device is connected to a power controller and wherein the signal from the control device is sent to the power controller that changes the power state of the first network device from unpowered to powered.

14. (Withdrawn) The method as recited in Claim 12, wherein the second network device is a terminal server and wherein identifying the alteration at the terminal server includes determining whether a state of a port of the terminal server is changed from dead to

active in response to changing the power state of the first network device from unpowered to powered.

15. (Withdrawn) The method as recited in Claim 12, wherein the second network device is a switch and wherein identifying the alteration at the switch includes determining whether a trap on a port of the switch is raised in response to changing the power state of the first network device from unpowered to powered.

16. (Withdrawn) A method for determining how devices are interconnected in a network, the method comprising the computer-implemented steps of:

power cycling a first network device from either "off" to "on" or from "on" to "off";

identifying whether a suspected link of the first network device and a second network device becomes active as a result of power cycling of the first network device; and

when the suspected link become active, creating and storing information representing that the first network device is connected to the second network device.

17.-19. (Cancelled)

20. (Previously Presented) A computer-readable storage medium storing one or more sequences of instructions for determining one or more logical interconnections among a plurality of network devices that are interconnected in a network in an indefinite relationship, wherein a power state is associated with a first network device, which

instructions, when executed by one or more processors, cause the one or more processors to carry out the steps of:

changing the power state of the first network device from either (a) an unpowered state to a powered state or (b) from the powered state to the unpowered state;

identifying whether an alteration occurs at a second network device in response to changing the power state of the first network device; and

when the alteration occurs at the second network device, creating and storing first information representing a logical connection of the first network device to the second network device.

21. (Original) The computer-readable storage medium as recited in Claim 20, further comprising instructions which, when executed by one or more processors, cause the one or more processors to carry out the steps of:

retrieving second information from a database, wherein the second information represents one or more logical connections of the first network device to the second network device;

comparing the second information from the database with the first information; and

generating an error if the second information indicates that a logical connection exists between the first and second network devices but the first information does not indicate that the logical connection exists between the first and second network devices.

22. (Original) The computer-readable storage medium as recited in Claim 20, wherein the second network device is a terminal server and wherein the step of identifying whether

the alteration occurs at the terminal server further comprises instructions which, when executed by one or more processors, cause the one or more processors to carry out the step of:

determining whether a state of a port of the terminal server is changed from dead to active in response to changing the power state of the first network device.

23. (Original) The computer-readable storage medium as recited in Claim 20, wherein the second network device is a switch and wherein the step of identifying whether the alteration occurs at the switch further comprises instructions which, when executed by one or more processors, cause the one or more processors to carry out the step of:

determining whether a trap on a port of the switch is raised in response to changing the power state of the first network device.

24. (Original) The computer-readable storage medium as recited in Claim 20, further comprising instructions which, when executed by one or more processors, cause the one or more processors to carry out the steps of:

receiving, in response to changing the power state of the first network device, additional information from the first network device; and

recording the additional information.

25. (Original) The computer-readable storage medium as recited in Claim 20, wherein changing the power state of the first network device is in response to a signal from a third network device.

26. (Original) The computer-readable storage medium as recited in Claim 25, wherein the first network device is connected to a power controller and wherein the signal from the third network device is sent to the power controller that changes the power state of the first network device.

27. (Withdrawn) A computer-readable storage medium storing one or more sequences of instructions for determining one or more logical interconnections among a set of specified network devices that are interconnected in a network in an indefinite relationship, which instructions, when executed by one or more processors, cause the one or more processors to carry out the steps of:

(1) establishing connections among a plurality of network devices based upon a set of rules;

(2) activating a particular network device of said set of specified network devices by supplying power to the particular network device that previously was not supplied with power;

(3) identifying whether, in response to activating the particular network device, a change occurs at one or more network devices of said plurality of network devices;

(4) when the change occurs at each of the one or more network devices, creating and storing information representing a logical connection of the particular network device to each of the one or more network devices; and

(5) repeating steps (2), (3), and (4) for each of said set of specified network devices.

28. (Withdrawn) A computer-readable medium carrying one or more sequences of instructions for determining how devices are interconnected in a network, which

instructions, when executed by one or more processors, cause the one or more processors to carry out the steps of:

sending a signal from a control device that results in a change in a power state of a first network device in response to the signal, wherein the power state changes from either powered to unpowered or from unpowered to powered;

determining whether the first network device is connected to a second network device by identifying an alteration at the second network device that occurs in response to changing the power state of the first network device; and

when the alteration occurs at the second network device, creating and storing information representing that the first network device is connected to the second network device.

29. (Withdrawn) A computer-readable storage medium carrying one or more sequences of instructions for determining how devices are interconnected in a network, which instructions, when executed by one or more processors, cause the one or more processors to carry out the steps of:

power cycling a first network device from either "off" to "on" or from "on" to "off";

identifying whether a suspected link of the first network device and a second network device becomes active as a result of power cycling of the first network device; and

when the suspected link become active, creating and storing information representing that the first network device is connected to the second network device.

30. (Cancelled)

31. (Previously Presented) An apparatus for determining one or more logical interconnections among a plurality of network devices that are interconnected in a network in an indefinite relationship, wherein a power state is associated with a first network device, the apparatus comprising:

a means for changing the power state of the first network device from either (a) an unpowered state to a powered state or (b) from the powered state to the unpowered state;

a means for identifying whether an alteration occurs at a second network device in response to changing the power state of the first network device; and

a means for creating and storing first information representing a logical connection of the first network device to the second network device, when the alteration occurs at the second network device.

32. (Original) The apparatus as recited in Claim 31, further comprising:

a means for retrieving second information from a database, wherein the second information represents one or more logical connections of the first network device to the second network device;

a means for comparing the second information from the database with the first information; and

a means for generating an error if the second information indicates that a logical connection exists between the first and second network devices but the first information does not indicate that the logical connection exists between the first and second network devices.

33. (Original) The apparatus as recited in Claim 31, wherein the second network device is a terminal server and wherein the means for identifying whether the alteration occurs at the terminal server further comprises:

a means for determining whether a state of a port of the terminal server is changed from dead to active in response to changing the power state of the first network device.

34. (Original) The apparatus as recited in Claim 31, wherein the second network device is a switch and wherein the means for of identifying whether the alteration occurs at the switch further comprises:

a means for determining whether a trap on a port of the switch is raised in response to changing the power state of the first network device.

35. (Original) The apparatus as recited in Claim 31, further comprising:

a means for receiving, in response to changing the power state of the first network device, additional information from the first network device; and

a means for recording the additional information.

36. (Original) The apparatus as recited in Claim 31, wherein changing the power state of the first network device is in response to a signal from a third network device.

37. (Original) The apparatus as recited in Claim 36, wherein the first network device is connected to a power controller and wherein the signal from the third network device is sent to the power controller that changes the power state of the first network device.

38. (Withdrawn) An apparatus for determining one or more logical interconnections among a set of specified network devices that are interconnected in a network in an indefinite relationship, the apparatus comprising:

(1) a means for establishing connections among a plurality of network devices based upon a set of rules;

(2) a means for activating a particular network device of said set of specified network devices by supplying power to the particular network device that previously was not supplied with power;

(3) a means for identifying whether, in response to activating the particular network device, a change occurs at one or more network devices of said plurality of network devices;

(4) a means for creating and storing information representing a logical connection of the particular network device to each of the one or more network devices, when the change occurs at each of the one or more network devices; and

(5) a means for repeating steps (2), (3), and (4) for each of said set of specified network devices.

39. (Withdrawn) An apparatus for determining how devices are interconnected in a network, the apparatus comprising:

a means for sending a signal from a control device that results in a change in a power state of a first network device in response to the signal, wherein the power state changes from either powered to unpowered or from unpowered to powered;

a means for determining whether the first network device is connected to a second network device by identifying an alteration at the second network device that occurs in response to changing the power state of the first network device; and

a means for creating and storing information representing that the first network device is connected to the second network device, when the alteration occurs at the second network device.

40. (Withdrawn) An apparatus for determining how devices are interconnected in a network, the apparatus comprising:

a means for power cycling a first network device from either “off” to “on” or from “on” to “off”;

a means for identifying whether a suspected link of the first network device and a second network device becomes active as a result of power cycling of the first network device; and

a means for creating and storing information representing that the first network device is connected to the second network device, when the suspected link become active.

41. (Cancelled)

42. (Withdrawn) The computer-readable storage medium as recited in Claim 27, wherein the set of rules are applied based upon one or more attributes of each connection.

43. (Withdrawn) The computer-readable storage medium as recited in Claim 42, wherein the one or more attributes of each connection include information that is selected from the group consisting of a type of connection between two or more network devices, the number of connections between a specific network device and one or more other network devices, and that a particular connection is between a first type of network device and a second type of network device.

44.-45. (Cancelled)

46. (Withdrawn) The computer-readable storage medium as recited in Claim 27, wherein the instruction for identifying whether the change occurs at one or more network devices further comprises instructions which, when executed by one or more processors, cause the one or more processors to carry out the step of:

determining whether a trap on a port of each of the one or more network devices is raised in response to activating the particular network device by supplying power to the particular network device that previously was not supplied with power.

47. (Withdrawn) The computer-readable storage medium as recited in Claim 28 wherein the first network device is connected to a power controller and wherein the signal from the control device is sent to the power controller that changes the power state of the first network device from unpowered to powered.

48. (Withdrawn) The computer-readable storage medium as recited in Claim 28, wherein the second network device is a terminal server and wherein identifying the alteration at the terminal server includes determining whether a state of a port of the terminal server is changed from dead to active in response to changing the power state of the first network device from unpowered to powered.

49. (Withdrawn) The computer-readable storage medium as recited in Claim 28, wherein the second network device is a switch and wherein identifying the alteration at the switch includes determining whether a trap on a port of the switch is raised in response to changing the power state of the first network device from unpowered to powered.

50. (Withdrawn) The apparatus as recited in Claim 38, wherein the set of rules are applied based upon one or more attributes of each connection.

51. (Withdrawn) The apparatus as recited in Claim 50, wherein the one or more attributes of each connection include information that is selected from the group consisting of a type of connection between two or more network devices, the number of connections between a specific network device and one or more other network devices, and that a particular connection is between a first type of network device and a second type of network device.

52.-53. (Cancelled)

54. (Withdrawn) The apparatus as recited in Claim 38, wherein the means for identifying whether the change occurs at one or more network devices further comprises:

means for determining whether a trap on a port of each of the one or more network devices is raised in response to activating the particular network device by supplying power to the particular network device that previously was not supplied with power.

55. (Withdrawn) The apparatus as recited in Claim 39 wherein the first network device is connected to a power controller and wherein the signal from the control device is sent to the power controller that changes the power state of the first network device from unpowered to powered.

56. (Withdrawn) The apparatus as recited in Claim 39, wherein the second network device is a terminal server and wherein identifying the alteration at the terminal server includes determining whether a state of a port of the terminal server is changed from dead to active in response to changing the power state of the first network device from unpowered to powered.

57. (Withdrawn) The apparatus as recited in Claim 39, wherein the second network device is a switch and wherein identifying the alteration at the switch includes determining whether a trap on a port of the switch is raised in response to changing the power state of the first network device from unpowered to powered.

58. (Previously Presented) The method as recited in Claim 1, wherein changing the power state of the first network device is in response to a signal from a third network device.

59. (Previously Presented) The method as recited in Claim 58, wherein the first network device is connected to a power controller and wherein the signal from the third network device is sent to the power controller that changes the power state of the first network device.

60. (Previously Presented) The method as recited in Claim 1, wherein:

when the power state of the first network device is the unpowered state, the first network device is not able to receive one or more packets over the network; and

when the power state of the first network device is the powered state, the first network device is able to receive one or more packets over the network.

61. (Previously Presented) The computer-readable storage medium as recited in Claim 20, wherein:

when the power state of the first network device is the unpowered state, the first network device is not able to receive one or more packets over the network; and

when the power state of the first network device is the powered state, the first network device is able to receive one or more packets over the network.

62. (Previously Presented) The apparatus as recited in Claim 31, wherein:

when the power state of the first network device is the unpowered state, the first network device is not able to receive one or more packets over the network; and

when the power state of the first network device is the powered state, the first network device is able to receive one or more packets over the network.

63. (Withdrawn) The method as recited in Claim 6, wherein:

when the particular network device is not supplied with power, the particular network device is not able to receive one or more packets over the network; and

when the particular network device is supplied with power, the particular network device is able to receive one or more packets over the network.

64. (Withdrawn) The computer-readable storage medium as recited in Claim 27, wherein:

when the particular network device is not supplied with power, the particular network device is not able to receive one or more packets over the network; and

when the particular network device is supplied with power, the particular network device is able to receive one or more packets over the network.

65. (Withdrawn) The apparatus as recited in Claim 38, wherein:

when the particular network device is not supplied with power, the particular network device is not able to receive one or more packets over the network; and

when the particular network device is supplied with power, the particular network device is able to receive one or more packets over the network.

66. (Withdrawn) The method as recited in Claim 12, wherein:

when the power state of the first network device is unpowered, the first network device is not able to receive one or more packets over the network; and

when the power state of the first network device is powered, the first network device is able to receive one or more packets over the network.

67. (Withdrawn) The computer-readable storage medium as recited in Claim 28, wherein:

when the power state of the first network device is unpowered, the first network device is not able to receive one or more packets over the network; and

when the power state of the first network device is powered, the first network device is able to receive one or more packets over the network.

68. (Withdrawn) The apparatus as recited in Claim 39, wherein:

when the power state of the first network device is unpowered, the first network device is not able to receive one or more packets over the network; and

when the power state of the first network device is powered, the first network device is able to receive one or more packets over the network.

IX. EVIDENCE APPENDIX PAGE

None.

X. RELATED PROCEEDINGS APPENDIX PAGE

None.